

UNITED STATES UTILITY PATENT APPLICATION

FOR

A Method, System, and Apparatus for Dynamically

5

Configuring the Operating Point Utilized for

Thermal Management of an Integrated Circuit

Inventors:

10

Varghese George
Stephen H. Gunther
Sanjeev Jahagirdar
Inder Sodhi

15

Prepared by:
Michael Nesheiwat
Patent Attorney

20



Intel Corporation
2111 N.E. 25th Avenue;
JF3-147
Hillsboro, OR 97124
Phone: (503) 712-8918
Facsimile: (503) 264-1729

25

30

Express Mail Number – EV325527961US

BACKGROUND

[001] 1. Field

[002] The present disclosure pertains to the field of thermal and power management.

More particularly, the present disclosure pertains to dynamically configuring the

5 operating point as utilized for thermal management of an integrated circuit (IC).

[003] 2. Description of Related Art

[004] Power management schemes allow for dynamically adjusting the power

consumption of an electronic component in order to meet the power consumption

requirements for various types of systems and integrated devices, such as, servers,

10 laptops, integrated devices and desktops. Typically, thermal management schemes

provide a means for adjusting power consumption in order to limit the integrated device's

operating temperature. The maximum allowable operating temperature for an integrated

circuit is a function of the specific characteristics of a given unit. For example, units that

have a higher operating frequency may be allowed to operate at higher temperatures than

15 other units of the same design that have lower operating frequencies. The allowable

operating temperature of a device may also be a function of the voltage supplied to the

device or some other characteristic that varies across the distribution of parts built from a

single design. Because these characteristics may change over the multi-year lifetime

during which a specific design is manufactured, the allowable operating temperatures

20 may also change over this manufacturing lifetime. Likewise, different units of the same

design will have different power consumption. As a result, the power reduction necessary

to meet a thermal management requirement may change from unit to unit across the

manufacturing lifetime of the design. Hence, taking an identical action in response to a

thermal event on all parts manufactured from a particular design may be non-optimal.

[0005] One means of configuring the operating point for a thermal management event would be to allow software programmability. With this approach, software must know the relevant operating characteristics of that specific part, However, such a scheme may be
5 susceptible to software tampering to allow improper and unsupported voltage, frequency, or thermal operating points.

Brief Description of the Figures

[006] The present invention is illustrated by way of example and not limitation in the Figures of the accompanying drawings.

5 [007] Figure 1 illustrates a chart for a method utilized in accordance with an embodiment

[008] Figure 2 illustrates a schematic utilized in accordance with an embodiment.

[009] Figure 3 illustrates a flowchart for a method utilized in accordance with an embodiment.

10 [010] Figure 4 illustrates a system in accordance with one embodiment.

15

20

Detailed Description

[011] The following description provides a method and apparatus for improved configuration of the operating point to be used in the event of a thermal management issue. In the following description, numerous specific details are set forth in order to 5 provide a more thorough understanding of the present invention. It will be appreciated, however, by one skilled in the art that the invention may be practiced without such specific details. Those of ordinary skill in the art, with the included descriptions, will be able to implement appropriate logic circuits without undue experimentation.

[012] As previously described, the present methods incorporate software to configure 10 the operating point utilized for dealing with a thermal issue on an integrated circuit. In contrast, the claimed subject matter is not dependent on software to define the operating points within an operating range for a processor. The claimed subject matter facilitates dynamically configuring an operating point used in response to a predetermined condition. For example, one predetermined condition may be a thermal issue as indicated 15 by a thermal sensor indicating the integrated device's temperature meets or exceeds a temperature threshold. Another example of a predetermined condition may be a power issue, wherein the integrated device or a system with at least one integrated device meets or exceeds a power consumption threshold. Consequently, the claimed subject matter responds to the predetermined condition by adjusting the frequency, such as, bus ratio, 20 and adjusting the voltage. This mechanism chooses the operating point based on the characteristics of that specific unit. Furthermore, the bus ratio to be used in response to the predetermined condition will be discussed in further detail in the following figures.

[013] Figure 1 illustrates a chart for a method utilized in accordance with an

embodiment. In one aspect, the claimed subject matter adjusts the operating point used for responding to a thermal issue based on the normal operating frequency and voltage of that specific processor. Likewise, the operating point adjustment is made upon detection of the predetermined condition, such as, for example, a thermal issue or power issue.

5 [014] The claimed subject matter facilitates a thermal monitor mechanism to allow one skilled in the art, such as, a system designer to reduce the cost of thermal solutions without compromising system integrity or reliability. For example, in one embodiment, frequency and voltage scaling are used for reducing power to maintain the integrated device's temperature within a predetermined operating range. In one embodiment, a
10 thermal sensor, coupled to the integrated device, indicates whether the integrated device's temperature meets or exceeds a predetermined threshold. If so, the thermal monitor is enabled to reduce the frequency and voltage. In the event the integrated device's temperature is less than the threshold, the thermal monitor is disabled and the integrated device operates at the original frequency and voltage before the thermal monitor was
15 enabled. For example, in one embodiment, the integrated device, such as, a processor, is operating at a normal operating point with a voltage of $V_{nominal}$. In the event of a thermal issue, such as, when the thermal sensor indicates the temperature of the integrated device meets or exceeds a predetermined threshold, the thermal monitor is enabled and the processor's logic facilitates a change in voltage and frequency in response to the thermal
20 issue to reduce power consumption. Consequently, the frequency is changed to a value within an desired operating range, which is discussed further in connection with Figures 1-4 below, and the voltage is reduced to the minimum level that still supports operation at the new frequency.

[015] The chart depicts a two dimensional graph, an x-axis 104 and a y-axis 102 that depict an operating voltage of the processor and a bus ratio of the processor, respectively. In one embodiment, the bus ratio is defined as the processor's core frequency divided by a bus frequency. In the same embodiment, the chart depicts 5 defining an operating range 106 when the thermal monitor mechanism is enabled, such as, in one embodiment, in the event of a thermal issue when the thermal sensor indicates the temperature of the integrated device meets or exceeds a predetermined threshold . Therefore, the chart depicts an operating range 106 at a predetermined first voltage, for the processor when the thermal monitor is enabled; otherwise, the 10 processor is allowed to operate in a range defined as between 0 and Vmax for the voltage and Min Ratio to a Maximum Bus Ratio (Max Ratio). In the same embodiment, the Maximum Bus Ratio and the predetermined first voltage are determined during manufacturing test. Likewise, the predetermined first voltage for the operating range 106 is stored in a plurality of fuses. A ratio offset 108 will be 15 described in further detail in connection with Figure 2.

[016] In one embodiment, the predetermined first voltage is determined during silicon testing. Likewise, in one embodiment, the predetermined first voltage is the lowest voltage that is sufficient for the processor to perform at the frequency (bus ratio) chosen in response to the predetermined condition . Furthermore, 20 the the predetermined first voltage is programmed into an integrated device, for example, by storing it into a plurality of fuses.

[017] In one embodiment, the Max Ratio is the highest ratio that an integrated device, such as, a processor may perform without errors. Likewise, it is programmed into fuses.

In one embodiment, the fuses are programmed with the Max Ratio at production testing.

In contrast, in one embodiment, the startup bus ratio is not necessarily the same as the Max Ratio. For example, an integrated device for a mobile product may utilize the Min Ratio for a boot up operation. In contrast, an integrated device for a desktop product, the startup ratio is the same as the Max Ratio.

[018] Therefore, the claimed subject matter facilitates a thermal monitor mechanism to minimize performance impact as the operating voltage and frequency are lowered to accommodate the thermal event. As described earlier, the claimed subject matter determines the the predetermined first voltage, which is the lowest voltage that is required for the processor to operate at the frequency utilized in response to a thermal issue.

Subsequently, the operating range is defined and one may select a frequency between Min Ratio and Max ratio based at least in part on the amount of power reduction that is needed. For example, if the thermal monitor is enabled, one may select an operating frequency that has the smallest amount of power reduction in response to the predetermined condition while having the best performance. In contrast, one may select an operating frequency of Min ratio that has the most amount of power reduction in response to the predetermined condition while having the worst performance.

[019] Figure 2 illustrates a schematic utilized in accordance with an embodiment. In

one embodiment, the schematic incorporates the logic for defining the operating range

used during response to the predetermined condition as depicted in connection with Figure 1. For example, a processor may support a plurality of different bus frequencies.

A bus ratio for the operating range 106 is determined by subtracting an offset from a

startup bus ratio, wherein the offset depends on the processor's bus frequency 204. For example, in one embodiment, the offset values are calculated for each of the supported bus frequencies and are determined based at least in part on power estimates and are coded into RTL as constants.

5

[020] Consequently, the offsets for all the supported bus frequencies are programmed into the CPU. In the event of a processor reset or startup, the multiplexer 205 selects one of the programmed offsets 202 based on the bus frequency 204. Subsequently, the offset 206 is subtracted from the startup bus ratio to determine an operating range bus ratio,

10 210, to be utilized by the processor when the thermal monitor is active. However, if operating range bus ratio is less than the minimum supported bus ratio, the thermal monitor mode is disabled.

[021] Figure 3 illustrates a flowchart for a method utilized in accordance with an

embodiment. In one embodiment, the flowchart 300 depicts defining a bus ratio to be

15 utilized by a processor upon activation of a thermal monitor. The flowchart depicts storing a plurality of offset values for a plurality of supported bus frequencies by the processor, as depicted in a block 302. Subsequently, one of the plurality of offset values is selected and is subtracted from a startup bus ratio of the processor to determine a bus ratio to be utilized by the processor when the thermal monitor is active, as depicted in a

20 block 304. In one embodiment, one of the plurality of offset values is selected based at least in part on a bus frequency of the processor. Consequently, comparing the bus ratio to be utilized by the processor if the thermal monitor is active to a minimum bus ratio that is supported by the processor, as depicted by a block 306. If the bus ratio is less than the

minimum bus ratio, then the thermal monitor is disabled and the processor operates at the voltage and frequency that were being utilized before entering the thermal monitor.

[023] Figure 4 depicts a system in accordance with one embodiment. The system in one embodiment is a processor 402 that is coupled to a chipset 404 that is coupled to a memory 406. For example, the chipset performs and facilitates various operations, such as, memory transactions between the processor and memory. In one embodiment, the system comprises one or all of the previous embodiments depicted in connection with Figures 1-3 of the specification to support a processor's thermal operating range upon enabling a thermal mode.

[024] While certain exemplary embodiments have been described and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on the broad invention, and that this invention not be limited to the specific constructions and arrangements shown and described, since various other modifications may occur to those ordinarily skilled in the art upon studying this

disclosure.